

# Understand EtherChannel Load Balance and Redundancy on Catalyst Switches

## Contents

---

### [Introduction](#)

### [Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Conventions](#)

### [Background Information](#)

### [Determine the Link Across Which to Send Traffic](#)

[Catalyst 6500/6000 Series](#)

[Cisco IOS](#)

[Catalyst 4500/4000 Series](#)

[Cisco IOS](#)

[Catalyst 2900XL/3500XL Series](#)

[Catalyst 3750/3560](#)

[Catalyst 2950/2955/3550](#)

[Catalyst 1900/2820](#)

[Catalyst 2948G-L3/4908G-L3](#)

[Matrix of Load Balancing Methods](#)

### [What Is PAgP and Where Do You Use It?](#)

### [ISL/802.1Q Trunking Support on EtherChannel](#)

### [Related Information](#)

---

## Introduction

This document describes the EtherChannel algorithm for load balance and redundancy on Cisco Catalyst switches.

## Prerequisites

### Requirements

There are no specific requirements for this document.

### Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

## Conventions

Refer to [Cisco Technical Tips Conventions](#) for more information on document conventions.

## Background Information

Fast EtherChannel allows multiple physical Fast Ethernet links to combine into one logical channel.

This allows load sharing of traffic among the links in the channel as well as redundancy in the event that one or more links in the channel fail.

Fast EtherChannel can be used to interconnect LAN switches, routers, servers, and clients via unshielded twisted pair (UTP) wiring or single-mode and multimode fiber.


This document refers to Fast EtherChannel, Gigabit EtherChannel, port channel, channel and port group with a single term, EtherChannel. The information in the document applies to all of these EtherChannels.

This document presents the concept of load balancing and redundancy on Cisco Catalyst switches with the use of the EtherChannel.

This document also covers the Port Aggregation Protocol (PAgP) and trunking support over EtherChannel.

This document does not cover how to configure EtherChannel on Catalyst switches.

Refer to the documents in the Related Information section for details on how to configure EtherChannel on Catalyst switches.



*Catalyst 6500, Catalyst 2950, EtherChannel*

## Determine the Link Across Which to Send Traffic

This section explains the load balance of packets over an EtherChannel group for each switch platform.

### Catalyst 6500/6000 Series

EtherChannel reduces part of the binary pattern that the addresses in the frame form to a numerical value

that selects one of the links in the channel in order to distribute frames across the links in a channel.

EtherChannel frame distribution uses a Cisco-proprietary hashing algorithm. The algorithm is deterministic; if you use the same addresses and session information, you always hash to the same port in the channel.

This method prevents out-of-order packet delivery.


The Cisco-proprietary hash algorithm computes a value in the range 0 to 7. With this value as a basis, a particular port in the EtherChannel is chosen.

The port setup includes a mask which indicates which values the port accepts for transmission. With the maximum number of ports in a single EtherChannel, which is eight ports, each port accepts only one value.

If you have four ports in the EtherChannel, each port accepts two values, and so forth. This table lists the ratios of the values that each port accepts, which depends on the number of ports in the EtherChannel:


Number of Ports in the EtherChannel	Load Balancing
8	1:1:1:1:1:1:1:1
7	2:1:1:1:1:1:1
6	2:2:1:1:1:1
5	2:2:2:1:1
4	2:2:2:2
3	3:3:2
2	4:4

---

 **Note:** This table only lists the number of values, which the hash algorithm calculates, that a particular port accepts. You cannot control the port that a particular flow uses. You can only influence the load balance with a frame distribution method that results in the greatest variety.


---

---

 **Note:** The hash algorithm cannot be configured or changed to load balance the traffic among the ports in an EtherChannel

---

---

 **Note:** The same Cisco-proprietary hash algorithm is also implemented in Cisco Catalyst 6500/6000 Series Switches that run Cisco IOS® software.

---

Therefore, you can only achieve perfect load balancing, even with random addresses, if you have two, four, or eight ports in the port channel.

## Cisco IOS

Catalyst 6500/6000 switches that run Cisco IOSsystem software support a maximum of 64 EtherChannels.

You can form an EtherChannel with up to eight compatibly configured LAN ports on any module in a Catalyst 6500/6000 series switch.

All LAN ports in each EtherChannel must be the same speed, and you must configure all the ports as either Layer 2 or Layer 3 LAN ports.

Catalyst 6500/6000 switches that run Cisco IOS system software uses a Cisco-proprietary hash algorithm.

EtherChannel load balancing can use MAC addresses, IP addresses, or Layer 4 port numbers with a Policy Feature Card 2 (PFC2) and either source mode, destination mode, or both.


The mode you select applies to all EtherChannels that you configure on the switch. Use the option that provides the greatest variety in your configuration.

For example, if the traffic on a channel only goes to a single MAC address, use of the destination MAC address results in the choice of the same link in the channel each time.

Use of source addresses or IP addresses can result in a better load balance.

Issue the **port-channel load-balance {src-mac | dst-mac | src-dst-mac | src-ip | dst-ip | src-dst-ip | src-port | dst-port | src-dst-port | mpls}** global configuration command in order to configure the load balancing.

---

 **Note:** The **port-channel load-balance src-dst-mixed-ip-port** command can change the hardware forwarding on PFC/DFC/CFC of Supervisor and can cause interruption to traffic for several seconds to minutes until the new hash algorithm has been calculated and starts to forward the traffic accordingly. Any change in the hash algorithm is recommended during non-production hours.

---

Issue the [show etherchannel load-balance](#) command in order to check the frame distribution policy.

You can determine which interface in the EtherChannel forwards traffic, with the frame distribution policy as a basis.

Issue the [remote login switch](#) command to log in remotely to the Switch Processor (SP) console in order to make this determination.

Then, issue the **test etherchannel load-balance interface port-channelnumber{ip | l4port | mac} [source\_ip\_add|source\_mac\_add|source\_l4\_port] [dest\_ip\_add|dest\_mac\_add|dest\_l4\_port]** command.

These are some examples:

1.

```
<#root>

6509#

remote login switch

    Trying Switch ...
    Entering CONSOLE for Switch
    Type "^C^C^C" to end this session

6509-sp#

test etherchannel load-balance interface port-channel 1 ip 10.10.10.2 10.10.10.1

Would select Gi6/1 of Po1

6509-sp#
```

2.

```
<#root>

6509#

remote login switch

    Trying Switch ...
    Entering CONSOLE for Switch
    Type "^C^C^C" to end this session

6509-sp#

test etherchannel load-balance interface port-channel 1 mac 00d0.c0d7.2dd4 0002.fc26.2494

Would select Gi6/1 of Po1

6509-sp#
```

## Catalyst 4500/4000 Series

### Cisco IOS

A Catalyst 4500/4000 series switch with Cisco IOS Software (Supervisor Engine II+ and later) supports a maximum of 64 EtherChannels.

You can form an EtherChannel with up to eight compatibly configured Ethernet interfaces on any module and across modules.

All interfaces in each EtherChannel must be the same speed, and you must configure all the interfaces as either Layer 2 or Layer 3 interfaces.

EtherChannel reduces part of the binary pattern that is formed from the addresses in the frame to a numerical value that selects one of the links in the channel in order to balance the traffic load across the links in a channel.

EtherChannel load balancing can use MAC addresses, IP addresses, or Layer 4 port numbers and either source mode, destination mode, or both.


Use the option that provides the greatest variety in your configuration. For example, if the traffic on a channel only goes to a single MAC address, use of the destination MAC address results in the choice of the same link in the channel each time.

Use of source or IP addresses can result in a better load balance.

Issue the **port-channel load-balance {src-mac | dst-mac | src-dst-mac | src-ip | dst-ip | src-dst-ip | src-port | dst-port | src-dst-port}** global configuration command in order to configure load balancing.

Load Balance must be configured globally and the load balancing option cannot be changed on a per port basis.

---

 **Note:** The switch uses the lower order bits of source MAC address and destination MAC address in order to determine which links must be used to transmit the data. So, if the data is received from the same source, then same link of the EtherChannel is used in order to forward the data.

---

## Catalyst 2900XL/3500XL Series

A Catalyst 2900XL that runs a Cisco IOS software release that is earlier than Cisco IOS Software Release 11.2(8)SA3 chooses a link in the channel based on the link on which the destination MAC address was last heard.

The software dynamically reallocates this address to another link in the channel if the link on which the address was learned is busier than the others.

You can configure a Catalyst 2900XL that runs Cisco IOS Software Release 11.2(8)SA3 or later and a Catalyst 3500XL that runs Cisco IOS Software Release 11.2(8)SA6 or later in order to choose a link to be sent across the Fast EtherChannel.

The switch chooses the link on the basis of the destination or source MAC address of the frame. The default is to use the source MAC address.

This default means that all packets that the switch receives on a non-Fast EtherChannel port with the same MAC source address that have a destination of the MAC addresses on the other side of the channel take the same link in the channel.

Use source-based forwarding when many stations that are attached to the Catalyst 2900XL/3500XL send to a few stations, such as a single router, on the other side of the Fast EtherChannel.

The use of source-based forwarding in this situation evenly distributes traffic across all links in the channel. Also, the Catalyst 2900XL/3500XL switches maintain a notion of a default port on which to transmit traffic, such as Spanning Tree Protocol (STP), multicasts, and unknown unicasts.

## Catalyst 3750/3560

The Catalyst 3750/3560 series switch can support up to eight compatibly configured Ethernet interfaces in an EtherChannel. The EtherChannel provides

full-duplex bandwidth up to 800 Mbps (Fast EtherChannel) or 8 Gbps (Gigabit EtherChannel) between your switch and another switch or host.

With Cisco IOS Software Release 12.2(20)SE and earlier, the number of EtherChannels has a limit of 12.

With Cisco IOS Software Release 12.2(25)SE and later, the number of EtherChannels has a limit of 48.

EtherChannel balances the traffic load across the links in a channel through the reduction of part of the binary pattern that the addresses in the frame form to a numerical value that selects one of the links in the channel.

EtherChannel load balancing can use MAC addresses or IP addresses, source or destination addresses, or both source and destination addresses. The mode applies to all EtherChannels that are configured on the switch.

You configure the load balancing and forwarding method with use of the **port-channel load-balance {dst-ip | dst-mac | src-dst-ip | src-dst-mac | src-ip | src-mac}** global configuration command.

You can find out which interface is used in the EtherChannel to forward traffic based on the load balancing method.

The command for this determination is **test etherchannel load-balance interface port-channel number {ip | mac} [source\_ip\_add|source\_mac\_add] [dest\_ip\_add|dest\_mac\_add]** .

## Catalyst 2950/2955/3550

The Catalyst 2950/2955 series switch can support up to eight compatibly configured Ethernet interfaces in an EtherChannel.

The EtherChannel can provide full-duplex bandwidth up to 800 Mbps (Fast EtherChannel) or 2 Gbps (Gigabit EtherChannel) between your switch and another switch or host.

The number of EtherChannels has the limit of six with eight ports per EtherChannel.

The Catalyst 3550 series switches support both Layer 2 and Layer 3 EtherChannel, with up to eight compatibly configured Ethernet interfaces.

The EtherChannel provides full-duplex bandwidth up to 800 Mbps (Fast EtherChannel) or 8 Gbps (Gigabit EtherChannel) between your switch and another switch or host.

The limit of the number of EtherChannels is the number of ports of the same type.

For the 2950/2955/3550 series switch, EtherChannel balances the traffic load across the links in a channel by random association with a newly learned MAC address with one of the links in the channel.

EtherChannel load balancing can use either source-MAC or destination-MAC address forwarding.

With source-MAC address forwarding, when packets are forwarded to an EtherChannel, the packets are distributed across the ports in the channel based on the source-MAC address of the incoming packet.

Therefore, to provide load balancing, packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel.

With destination-MAC address forwarding, when packets are forwarded to an EtherChannel, the packets are distributed across the ports in the channel based on the destination host MAC address of the incoming packet.


Therefore, packets to the same destination are forwarded over the same port, and packets to a different destination are sent on a different port in the channel.

For the 3550 series switch, when source-MAC address forwarding is used, load distribution based on the source and destination IP address is also enabled for routed IP traffic.

All routed IP traffic chooses a port based on the source and destination IP address. Packets between two IP hosts always use the same port in the channel, and traffic between any other pair of hosts can use a different port in the channel.

Issue the **port-channel load-balance {dst-mac | src-mac}** global configuration command in order to configure the load-balance and forward method.

---

 **Note:** The default port is used to transmit traffic, such as Spanning Tree Protocol (STP), multicasts, and unknown unicast. The default port can be identified from the output of the command [show etherchannel summary](#) by a notation of **d**.

---

## Catalyst 1900/2820

With the enablement of PAgP, the two possible methods of link determination are preserve order and maximize load balancing between the links on the Fast EtherChannel.

The [What Is PAgP and Where Do You Use It?](#) section of this document describes PAgP. The default is to maximize load balancing. PAgP is used to negotiate the configured method with the device at the other side of the channel.

If preserve order is configured, the device at the other side is instructed in order to use source-based transmissions so that the Catalyst 1900/2820 always receives packets with the same source MAC address on the same link in the channel.

This is the link that the Catalyst 1900/2820 always uses to send traffic to this MAC address.

If maximize load balancing is configured, PAgP tells the other side that it can distribute traffic arbitrarily, and unicast traffic is transmitted by the Catalyst 1900/2820 on the link where the source address was last seen.

This provides the maximum possible load-balancing configuration. When Fast EtherChannel is configured with PAgP disabled, the switch cannot negotiate with the partner about the switch learning capability.

Whether the switch preserves frame order depends on whether the Fast EtherChannel partner performs source-based distribution.

The Catalyst 1900/2820s also elects an active port. The active port is used for flooded traffic such as unknown unicast, unregistered multicast, and broadcast packets.

If the port-channel mode is on (PAgP disabled), the active port is the link with the highest priority value.

If the mode is desirable or auto (PAgP enabled), the active port is selected based on the priority of the links on the switch that has the higher Ethernet address.

When two ports on the switch with the higher Ethernet address have the same priority, the port with the lower ifIndex is selected.

## Catalyst 2948G-L3/4908G-L3

When one link fails, all traffic that previously used that link now uses the link next to it. For example, if Link 1 fails in a bundle, traffic that previously used Link 1 before the failure now uses Link 2.

## Matrix of Load Balancing Methods

This matrix consolidates the load balancing methods:

Platform	Address Used in XOR	Source-Based?	Destination-Based?	Source-Destination-Based?	Load Balancing Method—Configurable/Fixed?
6500/6000	Layer 2, Layer 3 addresses, Layer 4 information, or MPLS information <sup>2</sup>	Yes	Yes	Yes	Configurable
4500/4000	Layer 2, Layer 3 addresses, or Layer 4 information	Yes	Yes	Yes	Configurable



<b>2900XL/3500XL</b>	Layer 2 address only	Yes	Yes	—	Configurable
<b>3750/3560</b>	Layer 2 or Layer 3 address only	Yes	Yes	Yes	Configurable
<b>2950/2955/3550</b>	Layer 2 address only <sup>1</sup>	Yes	Yes	— <sup>1</sup>	Configurable
<b>1900/2820</b>	These platforms use a special method of load balancing. See the <a href="#">Catalyst 1900/2820</a> section for details.				

<sup>1</sup> For the 3550 series switch, when source-MAC address forwarding is used, load distribution based on the source and destination IP address is also enabled for routed IP traffic. All routed IP traffic chooses a port based on the source and destination IP address.

<sup>2</sup> For the 6500 series switches that run Cisco IOS, MPLS layer 2 information can also be used for load balancing MPLS packets.

## What Is PAgP and Where Do You Use It?

PAgP aids in the automatic creation of EtherChannel links. PAgP packets are sent between EtherChannel-capable ports in order to negotiate the formation of a channel. Some restrictions are deliberately introduced into PAgP. The restrictions are:

- PAgP does not form a bundle on ports that are configured for dynamic VLANs. PAgP requires that all ports in the channel belong to the same VLAN or are configured as trunk ports. When a bundle already exists and a VLAN of a port is modified, all ports in the bundle are modified to match that VLAN.
- PAgP does not group ports that operate at different speeds or port duplex. If speed and duplex change when a bundle exists, PAgP changes the port speed and duplex for all ports in the bundle.
- PAgP modes are off, auto, desirable, and on. Only the combinations auto-desirable, desirable-desirable, and on-on allow the formation of a channel. The device on the other side must have PAgP set to on if a device on one side of the channel does not support PAgP, such as a router.

These switches do not support PAgP:

- Catalyst 2900XL/3500XL
- Catalyst 2948G-L3/4908G-L3

## ISL/802.1Q Trunking Support on EtherChannel

You can configure EtherChannel connections with or without Inter-Switch Link Protocol (ISL)/IEEE 802.1Q trunking.

After the formation of a channel, the configuration of any port in the channel as a trunk applies the configuration to all ports in the channel. Identically configured trunk ports can be configured as an EtherChannel.

You must have all ISL or all 802.1Q; you cannot mix the two. ISL/802.1Q encapsulation, if enabled, takes place independently of the source/destination load-balancing mechanism of Fast EtherChannel.

The VLAN ID has no influence on the link that a packet takes. ISL/802.1Q simply enables that trunk to belong to multiple VLANs. If trunking is not enabled, all ports that are associated with the Fast

EtherChannel must belong to the same VLAN.

## **Related Information**

- [Cisco Technical Support & Downloads](#)